

Jiayuan Mao

EMAIL: jiayuanm@mit.com

WEBSITE: <https://jiayuanm.com>

RESEARCH INTEREST

I aim to build machines that can continually learn new knowledge from their experiences and reason across tasks, modalities, and environments : answer queries, infer human intentions, and make long-horizon plans spanning hours to days. As an AI scientist, I primarily use *robots* as my testbed. My work spans multiple fields in AI, including *robotics*, *machine learning*, *computer vision*, and *natural language processing*.

EDUCATION

2019-Present	Massachusetts Institute of Technology Ph.D. Student in the Department of Electrical Engineering and Computer Science Advisors : Joshua B. Tenenbaum and Leslie Pack Kaelbling. Ph.D. Thesis : <i>Learning, Reasoning, and Planning with Neuro-Symbolic Concepts</i>
2014-2019	Tsinghua University B.Eng. in Computer Science (Yao Class), Institute for Interdisciplinary Information Sciences Thesis : <i>Learning Sememe-based Dependency Structures</i>

AWARDS AND HONORS

2024	Best Paper, CoRL 2024 Workshop on Language and Robot Learning
2024	Rising Star in EECS
2024	Rising Star in Generative AI
2024	Best Undergraduate Student Paper, Annual Meeting of the Cognitive Science Society
2021	Qualcomm Fellowship Finalist
2020	Facebook Fellowship Finalist
2019	Best Paper Nomination, Annual Meeting of the Association for Computational Linguistics
2019	MIT Presidential Graduate Fellowship

PUBLICATION

46 peer-reviewed conference and journal publications. In total 53 publications.

Learning Linear Attention in Polynomial Time	ArXiv 2024
Morris Yau, Eykin Akyurek, Jiayuan Mao , Joshua B. Tenenbaum, Stefanie Jegelka, Jacob Andreas	
One-Shot Manipulation Strategy Learning by Making Contact Analogies	
CoRL 2024 Workshop on Learning Effective Abstractions for Planning	
Yuyao Liu*, Jiayuan Mao* , Joshua Tenenbaum, Tomás Lozano-Pérez, Leslie Pack Kaelbling	
Keypoint Abstraction using Large Models for Object-Relative Imitation Learning	
CoRL 2024 Workshop on Language and Robot Learning (Best Paper)	
Xiaolin Fang*, Bo-Ruei Huang*, Jiayuan Mao* , Jasmine Shone, Joshua B. Tenenbaum, Tomás Lozano-Pérez, Leslie Pack Kaelbling	
BLADE : Learning Compositional Behaviors from Demonstration and Language	CoRL 2024
WeiYu Liu*, Neil Nie*, Ruohan Zhang, Jiayuan Mao †, Jiajun Wu†	
Embodied Agent Interface : A Single Line to Evaluate LLMs for Embodied Decision Making	NeurIPS 2024 (Oral)
Manling Li*, Shiyu Zhao*, Qineng Wang*, Kangrui Wang*, Yu Zhou*, Sanjana Srivastava, Cem Gokmen, Tony Lee, Li Erran Li, Ruohan Zhang, WeiYu Liu, Percy Liang, Li Fei-Fei, Jiayuan Mao , Jiajun Wu	

Hybrid Declarative-Imperative Representations for Hybrid Discrete-Continuous Decision-Making	WAFR 2024
<u>Jiayuan Mao</u> , Joshua B. Tenenbaum, Tomás Lozano-Pérez, Leslie Pack Kaelbling	
Agent Workflow Memory	ArXiv 2024
Zora Zhiruo Wang, <u>Jiayuan Mao</u> , Daniel Fried, Graham Neubig	
What Makes a Maze Look Like a Maze ?	
ECCV Human-Inspired Computer Vision Workshop 2024	
Joy Hsu, <u>Jiayuan Mao</u> , Joshua B. Tenenbaum, Noah D. Goodman, Jiajun Wu	
Learning Iterative Reasoning through Energy Diffusion	ICML 2024
Yilun Du*, <u>Jiayuan Mao*</u> , Joshua B. Tenenbaum	
Finding Structure in Logographic Writing with Library Learning	CogSci 2024
Guangyuan Jiang, Matthias Hofer, <u>Jiayuan Mao</u> , Lionel Wong, Joshua B. Tenenbaum, Roger P. Levy	(Best Undergraduate Student Paper)
“Set It Up!” : Functional Object Arrangement with Compositional Generative Models	RSS 2024
Yiqing Xu, <u>Jiayuan Mao</u> , Yilun Du, Tomas Lozano-Pérez, Leslie Pack Kaelbling, David Hsu	
Grounding Language Plans in Demonstrations through Counter-Factual Perturbations	ICLR 2024 (Spotlight)
Yanwei Wang, Tsun-Hsuan Wang, <u>Jiayuan Mao</u> , Michael Hagenow, Julie Shah	
Learning to Act from Actionless Videos through Dense Correspondences	ICLR 2024 (Spotlight)
Po-Chen Ko, <u>Jiayuan Mao</u> , Yilun Du, Shao-Hua Sun, Joshua B. Tenenbaum	
Learning Adaptive Planning Representations with Natural Language Guidance	ICLR 2024
Lionel Wong*, <u>Jiayuan Mao*</u> , Pratyusha Sharma*, Zachary S. Siegel, Jiahai Feng, Noa Korneev, Joshua B. Tenenbaum, Jacob Andreas	
Learning Planning Abstractions from Language	ICLR 2024
Weiyu Liu, Geng Chen, Joy Hsu, Jiajun Wu*, <u>Jiayuan Mao*</u>	
What Planning Problem Can A Relational Neural Network Solve	NeurIPS 2023 (Spotlight)
<u>Jiayuan Mao</u> , Tomás Lozano-Pérez, Joshua B. Tenenbaum, Leslie Pack Kaelbling	
What’s Left ? Concept Grounding with Logic-Enhanced Foundation Models	NeurIPS 2023
Joy Hsu*, <u>Jiayuan Mao*</u> , Joshua B. Tenenbaum, Jiajun Wu	
Learning Reusable Manipulation Strategies	CoRL 2023
<u>Jiayuan Mao</u> , Tomás Lozano-Pérez, Joshua B. Tenenbaum, Leslie Pack Kaelbling	
Compositional Diffusion-Based Continuous Constraint Solvers	CoRL 2023
Zhutian Yang, <u>Jiayuan Mao</u> , Yilun Du, Jiajun Wu, Joshua B. Tenenbaum, Tomás Lozano-Pérez, Leslie Pack Kaelbling	
Composable Part-Based Manipulation	CoRL 2023
Weiyu Liu, <u>Jiayuan Mao</u> , Joy Hsu, Tucker Hermans, Animesh Garg, Jiajun Wu	
NS3D : Neuro-Symbolic Grounding of 3D Objects and Relations	CVPR 2023
Joy Hsu, <u>Jiayuan Mao</u> , Jiajun Wu	
Programmatically Grounded, Compositionally Generalizable Robotic Manipulation	ICLR 2023 (Spotlight)
Renhao Wang*, <u>Jiayuan Mao*</u> , Joy Hsu, Hang Zhao, Jiajun Wu, Yang Gao	
Learning Rational Subgoals from Demonstrations and Instructions	AAAI 2023
Zhezhen Luo*, <u>Jiayuan Mao*</u> , Jiajun Wu, Tomás Lozano-Pérez, Joshua B. Tenenbaum, Leslie Pack Kaelbling	

DisCo : Improving Compositional Generalization in Visual Reasoning through Distribution Coverage	TMLR 2023
Joy Hsu, <u>Jiayuan Mao</u> , Jiajun Wu	
On the Expressiveness and Generalization of Hypergraph Neural Networks	LoG 2022
Zhezhen Luo, <u>Jiayuan Mao</u> , Joshua B. Tenenbaum, Leslie Pack Kaelbling	
Sparse and Local Hypergraph Reasoning Networks	LoG 2022
Guangxuan Xiao, Leslie Pack Kaelbling, Jiajun Wu, <u>Jiayuan Mao</u>	
PDSketch : Integrated Domain Programming, Learning, and Planning	NeurIPS 2022
<u>Jiayuan Mao</u> , Tomás Lozano-Pérez, Joshua B. Tenenbaum, Leslie Pack Kaelbling	
HandMeThat : Human-Robot Communication in Physical and Social Environments	NeurIPS 2022
Yanming Wan*, <u>Jiayuan Mao*</u> , Joshua B. Tenenbaum	
CLEVRER-Humans : Describing Physical and Causal Events the Human Way	NeurIPS 2022
<u>Jiayuan Mao*</u> , Xuelin Yang*, Xikun Zhang, Noah D. Goodman, Jiajun Wu	
IKEA-Manual : Seeing Shape Assembly Step by Step	NeurIPS 2022
Ruo Cheng Wang, Yunzhi Zhang, <u>Jiayuan Mao</u> , Ran Zhang, Chin-Yi Cheng, Jiajun Wu	
Translating a Visual LEGO Manual to a Machine-Executable Plan	ECCV 2022
Ruo Cheng Wang, Yunzhi Zhang, <u>Jiayuan Mao</u> , Chin-Yi Cheng, Jiajun Wu	
Programmatic Concept Learning for Human Motion Description and Synthesis	CVPR2022
Sumith Kulal*, <u>Jiayuan Mao*</u> , Alex Aiken†, Jiajun Wu†	
FALCON : Fast Visual Concept Learning by Integrating Images, Linguistic descriptions, and Conceptual Relations	ICLR 2022
Lingjie Mei*, <u>Jiayuan Mao*</u> , Ziqi Wang, Chuang Gan, Joshua B. Tenenbaum	
Grammar-Based Grounded Lexicon Learning	NeurIPS 2021
<u>Jiayuan Mao</u> , Haoyue Shi, Jiajun Wu, Roger P. Levy, Joshua B. Tenenbaum	
Temporal and Object Quantification Networks	IJCAI 2021
<u>Jiayuan Mao*</u> , Zhezhen Luo*, Chuang Gan, Joshua B. Tenenbaum, Jiajun Wu, Leslie Pack Kaelbling, Tomer D. Ullman	
Language-Mediated, Object-Centric Representation Learning	ACL 2021 (Findings)
Ruo Cheng Wang*, <u>Jiayuan Mao*</u> , Samuel J. Gershman, Jiajun Wu	
Hierarchical Motion Understanding via Motion Programs	CVPR 2021
Sumith Kulal*, <u>Jiayuan Mao*</u> , Alex Aiken, Jiajun Wu	
Grounding Physical Concepts of Objects and Events Through Dynamic Visual Reasoning	ICLR 2021
Zhenfang Chen, <u>Jiayuan Mao</u> , Jiajun Wu, Kwan-Yee K. Wong, Joshua B. Tenenbaum, Chuang Gan	
Object-Centric Diagnosis of Visual Reasoning	ArXiv 2020
Jianwei Yang, <u>Jiayuan Mao</u> , Jiajun Wu, Devi Parikh, David D. Cox, Joshua B. Tenenbaum, Chuang Gan	
Multi-Plane Program Induction with 3D Box Priors	NeurIPS 2020
Yikai Li*, <u>Jiayuan Mao*</u> , Xiuming Zhang, William T. Freeman, Joshua B. Tenenbaum, Noah Snavely, Jiajun Wu	
Perspective Plane Program Induction from a Single Image	CVPR 2020
Yikai Li*, <u>Jiayuan Mao*</u> , Xiuming Zhang, William T. Freeman, Joshua B. Tenenbaum, Jiajun Wu	
Visual Concept-Metaconcept Learning	NeurIPS 2019
Chi Han*, <u>Jiayuan Mao*</u> , Chuang Gan, Joshua B. Tenenbaum, Jiajun Wu	

Program-Guided Image Manipulators	ICCV 2019
<u>Jiayuan Mao*</u> , Xiuming Zhang*, Yikai Li, William T. Freeman, Joshua B. Tenenbaum, Jiajun Wu	
Visually Grounded Neural Syntax Acquisition	ACL 2019 (Best Paper Nominee)
Haoyue Shi*, <u>Jiayuan Mao*</u> , Kevin Gimpel, Karen Livescu	
Neurally-Guided Structure Inference	ICML 2019
Sidi Lu*, <u>Jiayuan Mao*</u> , Joshua B. Tenenbaum, Jiajun Wu	
The Neuro-Symbolic Concept Learner : Interpreting Scenes, Words, and Sentences from Natural Supervision	ICLR 2019 (Oral)
<u>Jiayuan Mao</u> , Chuang Gan, Pushmeet Kohli, Joshua B. Tenenbaum, Jiajun Wu	
Neural Logic Machines	ICLR 2019
Honghua Dong*, <u>Jiayuan Mao*</u> , Tian Lin, Chong Wang, Lihong Li, Denny Zhou	
Unified Visual-Semantic Embeddings : Bridging Vision and Language with Structured Meaning Representations	CVPR 2019 (Oral)
Hao Wu*, <u>Jiayuan Mao*</u> , Yufeng Zhang, Yuning Jiang, Lei Li, Wei-Ying Ma	
Neural Phrase-to-Phrase Machine Translation	ArXiv 2018
Jiangtao Feng, Lingpeng Kong, Po-Sen Huang, Chong Wang, Da Huang, <u>Jiayuan Mao</u> , Kan Qiao, Denny Zhou	
Acquisition of Localization Confidence for Accurate Object Detection	ECCV 2018 (Oral)
Borui Jiang*, Ruixuan Luo*, <u>Jiayuan Mao*</u> , Tete Xiao, Yuning Jiang	
Learning Visually-Grounded Semantics from Contrastive Adversarial Samples	COLING 2018
Haoyue Shi*, <u>Jiayuan Mao*</u> , Tete Xiao*, Yuning Jiang, Jian Sun	
Universal Agent for Disentangling Environments and Tasks	ICLR 2018
<u>Jiayuan Mao</u> , Honghua Dong, Joseph J. Lim	
What Can Help Pedestrian Detection ?	CVPR 2017
<u>Jiayuan Mao*</u> , Tete Xiao*, Yuning Jiang, Zhimin Cao	

INVITED TALKS

2024	Workshop on Rising Stars in Generative AI Title : <i>Learning, Reasoning and Planning with Neuro-Symbolic Concepts</i>
2024	Workshop on Visual Concepts at ECCV 2024 Title : <i>Representation and Computation Aspects of Visual Concepts</i>
2024	Bimanual Manipulation : On Kitchen Challenges workshop at ICRA 2024 Title : <i>Building General-Purpose Robots with Neuro-Symbolic Action Abstractions</i>
2024	Brown Robotics Talks at Brown University Title : <i>Compositional Action Representations</i>
2024	NSF Workshop on Hardware-Software Co-design for Neuro-Symbolic Computation Title : <i>Learning and Planning with Neuro-Symbolic Actions</i>
2024	The Manipulation Reading Group at the Robotics Institute at Carnegie Mellon University Title : <i>Building General-Purpose Robots with Integrated Learning and Planning</i>
2024	Coordinated Science Laboratory Student Conference (CSLSC 2024) at the University of Illinois at Urbana-Champaign Title : <i>Integrated Learning and Planning</i>
2023	Guest Lecture of Course “Reinforcement Learning” at National Taiwan University Title : <i>Integrated Robotic Programming, Learning and Planning</i>

2023	Workshop on Robot Representations For Scene Understanding, Reasoning and Planning at RSS 2023 Title : <i>Neuro-Symbolic Concepts for Robotic Manipulation</i>
2021	CLVR Lab at University of Southern California Title : <i>Neuro-Symbolic Frameworks for Visual Concept Learning and Language Acquisition</i>
2020	Tutorial on Neuro-Symbolic Reasoning and Program Synthesis Title : <i>Neuro-Symbolic Frameworks for Visual Concept Learning and Language Acquisition</i>
2020	Computational Cognitive Neuroscience Lab at Harvard University Title : <i>Neuro-Symbolic Frameworks for Visual Concept Learning and Language Acquisition</i>
2020	MIT Vision Seminar Title : <i>Neuro-Symbolic Frameworks for Visual Concept Learning and Language Acquisition</i>
2019	Workshop on Visually Grounded Interaction and Language at NeurIPS 2019 Title : <i>Neuro-Symbolic Frameworks for Visually Grounded Reasoning and Language Acquisition</i>

MENTORED STUDENTS

21 undergraduate and master's student mentees, 3 Ph.D. student mentees

Undergraduates and Master Students

2024	Bo-Ruei Huang
2024	Peiqi Liu
2024	Yuyao Liu
2023-2024	Neil Nie
2023 -2024	Zachary S. Siegel
2023-2024	Po-Chen Ko
2023-2024	Guangyuan Jiang
2023-2024	Zachary Zhang Next : Stripe
2023	Xingjian Bai Next : Ph.D. student at Massachusetts Institute of Technology
2023	Jiahai Feng Next : Ph.D. student at University of California, Berkeley
2023	Jung-Chun Liu Next : Ph.D. student at University of Michigan
2022	Renhao Wang Next : Ph.D. student at University of California, Berkeley
2022	Guangxuan Xiao Next : Ph.D. student at Massachusetts Institute of Technology
2021-2022	Yanming Wan Next : Ph.D. student at University of Washington
2021-2022	Xuelin Yang Next : Ph.D. student at University of California, Berkeley
2020-2022	Zhezheng Luo Next : Citadel
2020-2021	Ruocheng Wang Next : Ph.D. student at Stanford University
2020-2021	Lingjie Mei Next : Ph.D. student at Princeton University

2019-2021	Yikai Li Next : Ph.D. student at Stanford University
2020	Ruidong Wu Next : Ph.D. student at University of Illinois Urbana-Champaign
2019	Chi Han Next : Ph.D. student at University of Illinois Urbana-Champaign

Ph.D. Students

2023-2024	Weiyu Liu Postdoc at Stanford University. Co-mentored with Prof. Jiajun Wu
2022-2024	Joy Hsu Ph.D. student at Stanford University. Co-mentored with Prof. Jiajun Wu
2020-2022	Sumith Kulal Ph.D. student at Stanford University. Co-mentored with Prof. Jiajun Wu and Prof. Alex Aiken

TEACHING

2021 Fall	Teaching Assistant : Representation, Inference and Reasoning in AI (Graduate), MIT
2017 Spring	Teaching Assistant : Object-Oriented Programming (Undergraduate), Tsinghua University

PROFESSIONAL SERVICE : WORKSHOP AND TUTORIAL ORGANIZATION

AAAI 2025	Workshop on Planning in the Era of LLMs
CoRL 2024	Workshop on Learning Effective Abstractions for Planning
ECCV 2024	Workshop on Visual Concepts
NAACL 2021	Workshop on Visually Grounded Interaction and Language
CVPR 2020	Tutorial on Neuro-Symbolic Visual Reasoning and Program Synthesis

PROFESSIONAL SERVICE : CONFERENCE REFEREEING

2024-Present	International Conference on Automated Planning and Scheduling (ICAPS)
2024-Present	Conference on Language Models(CoLM)
2024-Present	Robotics : Science and Systems (RSS)
2023-Present	Association for the Advancement of Artificial Intelligence (AAAI)
2023-Present	IEEE International Conference on Robotics and Automation (ICRA)
2023-Present	IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
2023-Present	Association for Computational Linguistics Rolling Review (ACL Rolling Review)
2022-Present	European Conference on Computer Vision (ECCV)
2021-Present	International Conference on Computer Vision (ICCV)
2021-Present	International Conference on Machine Learning (ICML)
2021-Present	International Conference on Learning Representations (ICLR)
2020-Present	Conference on Neural Information Processing Systems (NeurIPS)
2019-Present	IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)

OTHER RESEARCH APPOINTMENTS

- 2018-2019** | **COCOSCI Group, Massachusetts Institute of Technology**
Visiting Student, Advisor : Joshua B. Tenenbaum
Publication : *The Neuro-Symbolic Concept Learner : Interpreting Scenes, Words, and Sentences from Natural Supervision*
- 2018** | **Google AI China Center**
Research Intern, Mentor : Denny Zhou, Chong Wang
Publication : *Neural Logic Machines*
Publication : *Neural Phrase-to-Phrase Machine Translation*
- 2017** | **CLVR Lab, University of Southern California**
Visiting Student, Advisor : Joseph J. Lim
Publication : *Universal Agent for Disentangling Environments and Tasks*